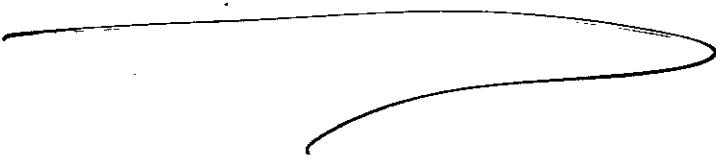


FNAL

27-28.3.00

CL2K



Are we $\gtrsim 95\%$ certain that we're
on the right track?

I. Incomplete summary

II. My 2¢

Glen Cowan

Royal Holloway College, Univ. London

28.3.00

My Paper

by Joe Author

I. Introduction

blah blah

II. Method

blah blah blah blah

III. Results

classical
statistics here

IV. Discussion & conclusions

Bayesian statistics
here

d'Agostini

Probability for Bayesians

Barlow's book

Very critical of Classical approach

e.g. What does limit on M_H mean?

Recently advocates L ratio for testing hypotheses

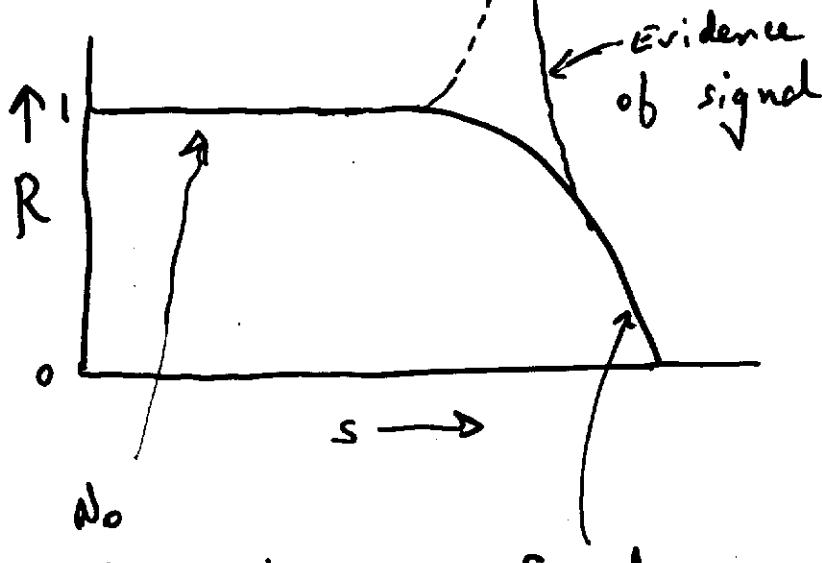
Used in Higgs limit [d'A + Degrassi]

$$\frac{P(s_1 | n_{\text{obs}}, b)}{P(s_2 | n_{\text{obs}}, b)} = \frac{P(n_{\text{obs}} | s_1, b)}{P(n_{\text{obs}} | s_2, b)} \times \frac{\text{Prior}(s_1)}{\text{Prior}(s_2)}$$

L ratio

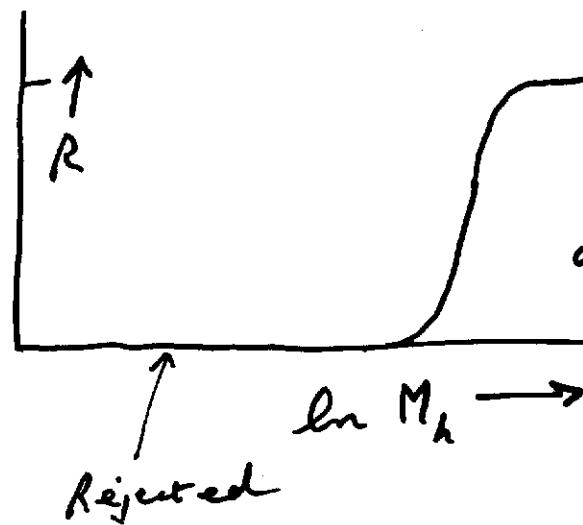
Rewrite as

$$\frac{P(s | n_{\text{obs}}, b)}{\text{Prior}(s)} / \frac{P(s=0 | n_{\text{obs}}, b)}{\text{Prior}(s=0)} = \frac{P(n_{\text{obs}} | s, b)}{P(n_{\text{obs}} | s=0, b)} = R$$



No discrimination

Signal rejected



Rejected

$R(s)$ is indep. of prior

From CERN workshop

(L. Lyons)

→ report likelihood function $\mathcal{L}(\vec{x}|\theta)$
data $\rightarrow \uparrow$
param

Or equivalent, e.g. $R = \frac{\mathcal{L}(\vec{x}|\theta)}{\mathcal{L}(\vec{x}|\theta=0)}$ (D'Agostis)
 \rightarrow plot

Why? So that a Bayesian consumer

can compute:

consumer's

✓ (subjective) prior

$$p(\theta|\vec{x}) \propto \mathcal{L}(\vec{x}|\theta)\pi(\theta)$$

Combining experiments: $\mathcal{L}(\theta) = \prod_i \mathcal{L}_i(\theta)$

$$\hookrightarrow \hat{\theta}_{ML}$$

- exact choice of model? (Bill Murray)
- systematics?
- $\mathcal{L}(\theta)$ not enough to compute Neyman conf. int. in small sample case.

The Four Commandments (Fred James)

- I Physicists should learn statistics vocab.
P-value, marginalize, . . .
- II State all assumptions, methods etc. in publication.
(and give an unbiased estimator $\hat{\theta}$, even if in unphysical region)
- III F-C in all searches
comments to follow . . .
- IV Bayesian decision theory in policy decision
and in discussion of results

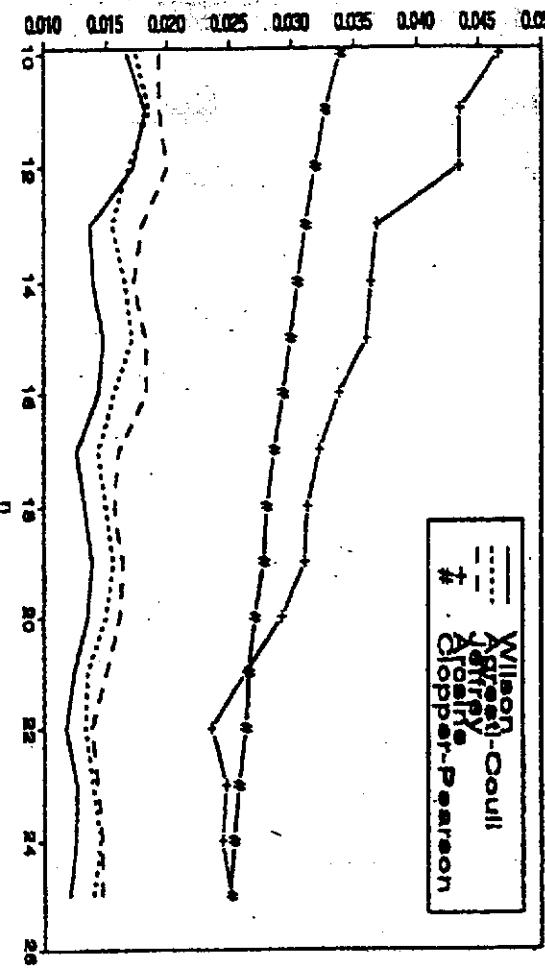
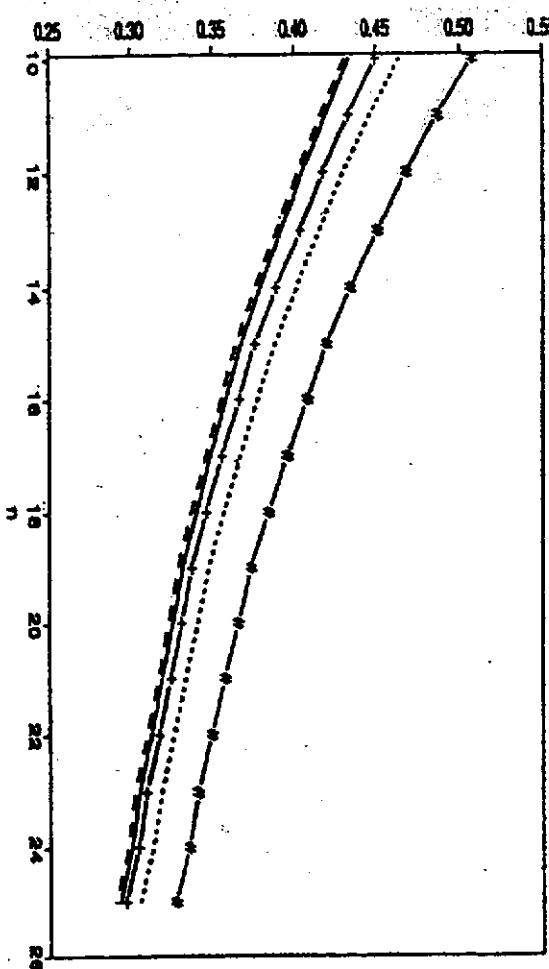
Subjective Bayes

$p(\theta | \vec{x})$ = degree of belief

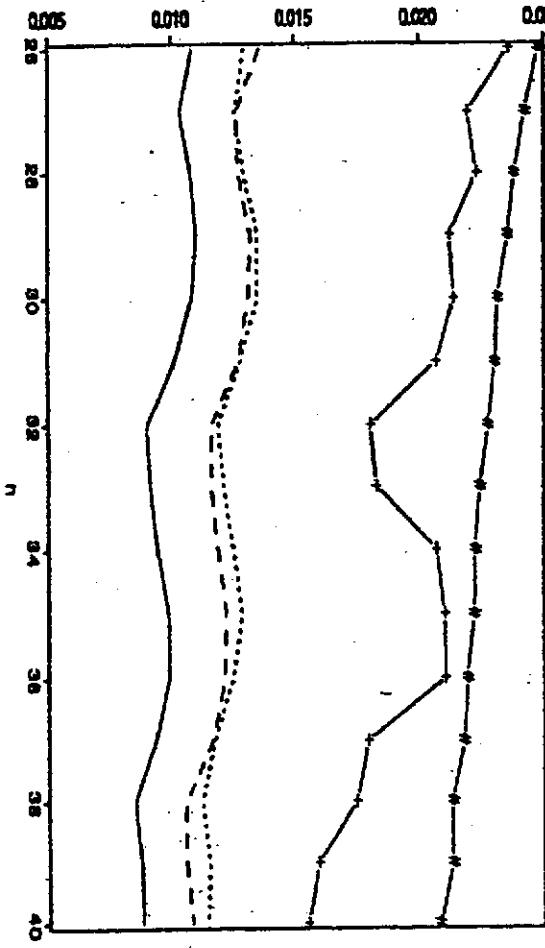
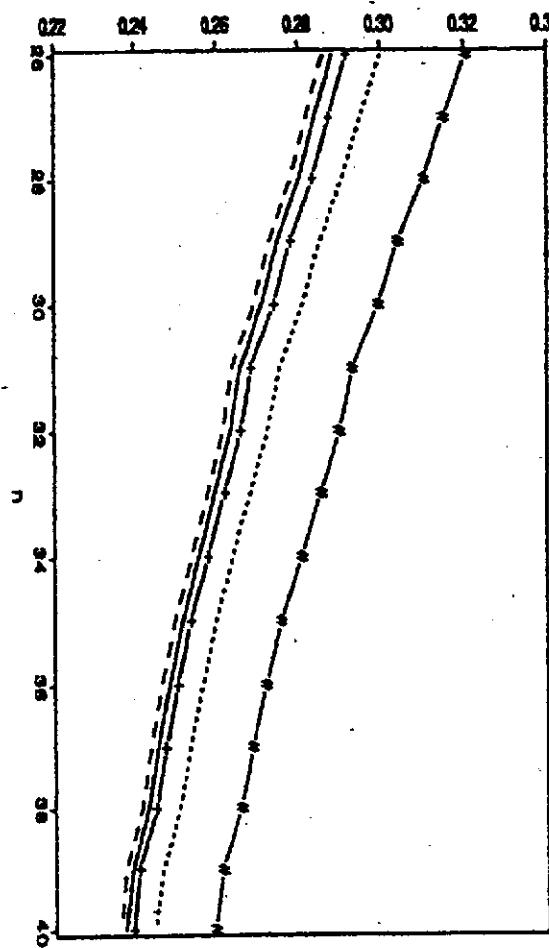
$$\propto \mathcal{L}(\vec{x} | \theta) \pi(\theta)$$

- By itself, a bad way to report result of an observation (Reindeer problem)
- Good for discussion / conclusions section of paper.
- It's how the consumer will digest the result, using his/her own prior.
(Explicitly or otherwise.)

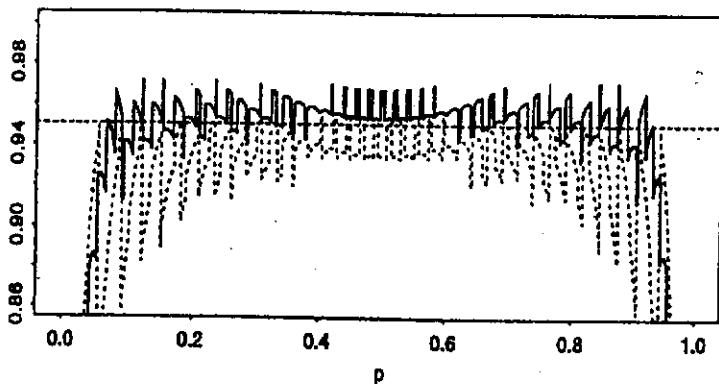
Mean Absolute Error



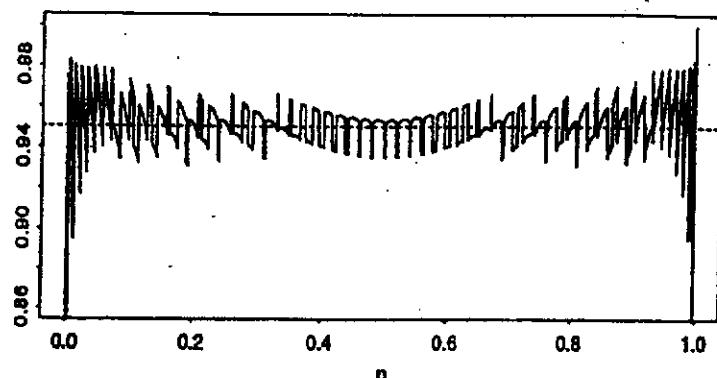
Mean Absolute Error



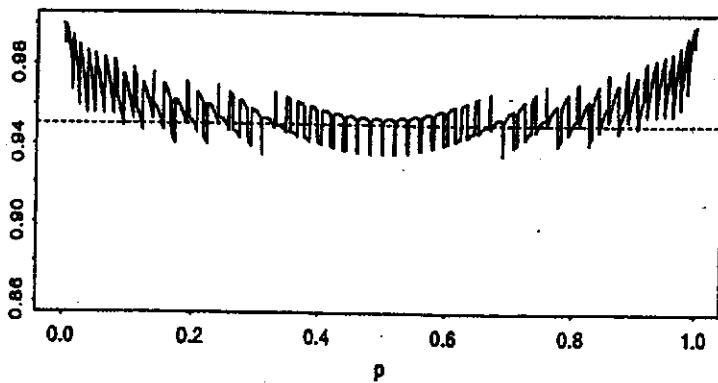
Recentered Interval



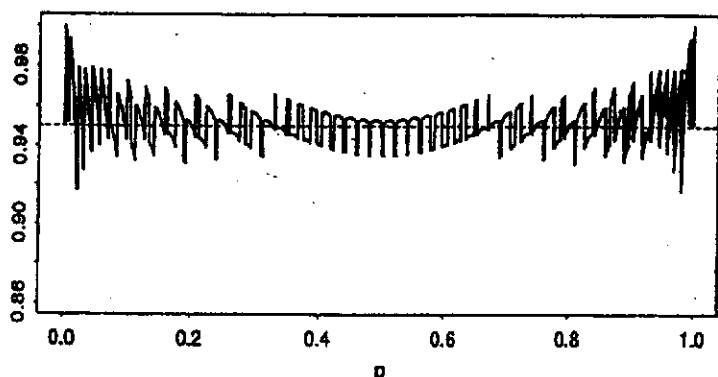
Wilson Interval



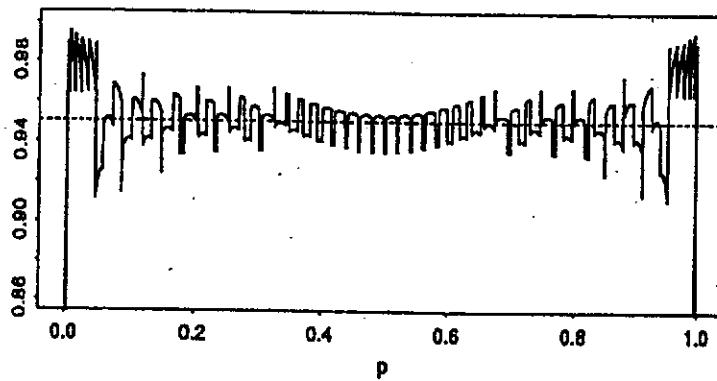
Agresti-Coull Interval



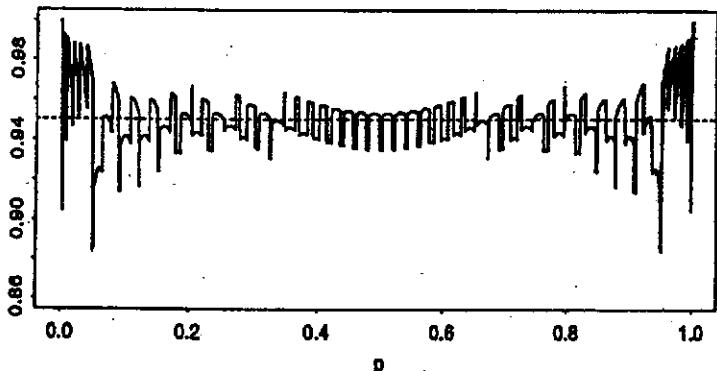
Modified Wilson Interval



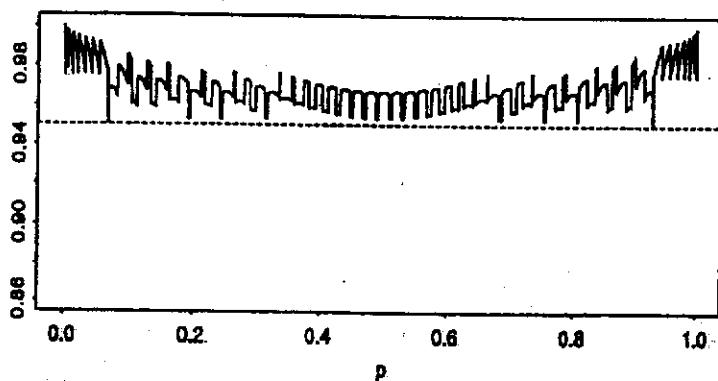
Arcsine Interval



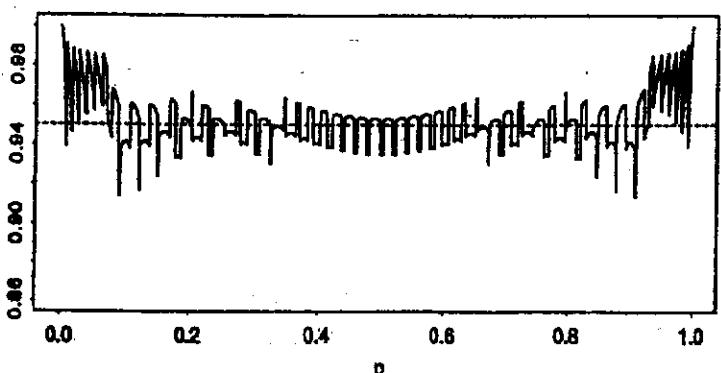
Jeffrey's Equal-Tailed Interval



Clopper-Pearson's Exact Interval



Modified Jeffrey's Equal-Tailed Interval



Empirical coverage of textbook confidence intervals

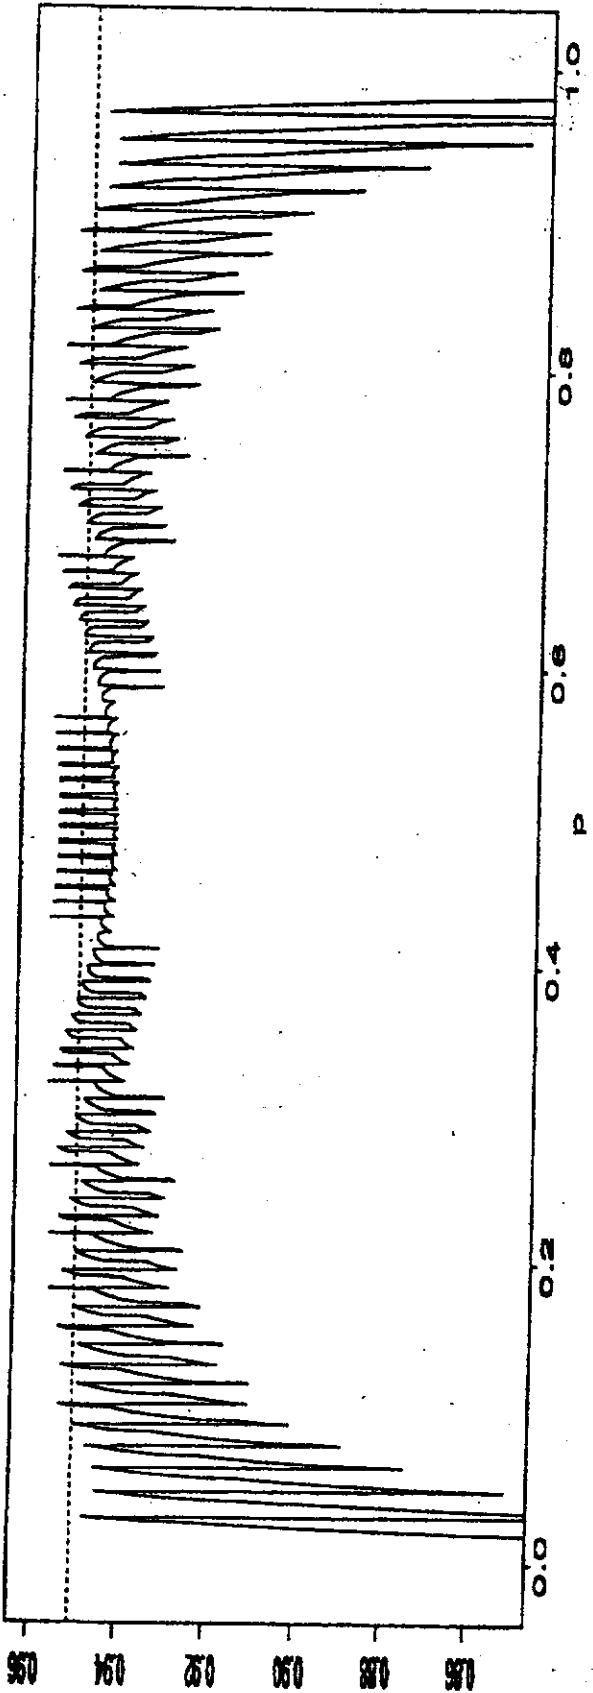


Figure 3: Oscillation phenomenon for fixed $n = 100$ and variable θ_0 ;
Nominal 95% Confidence interval

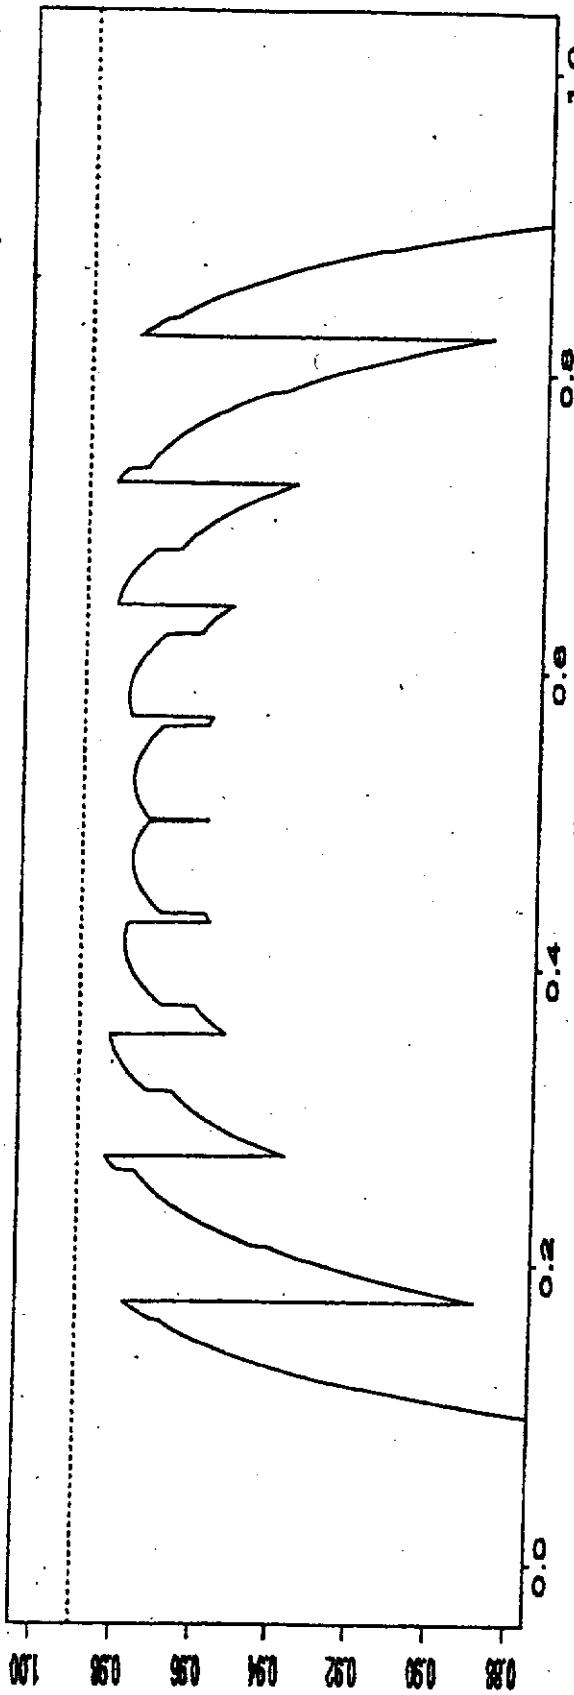


Figure : $n = 20$ and variable θ_0 ; nominal 99% interval.

"Objective Bayes"

- A prescription with Bayesian origins for constructing functions of the data $\theta_{np}(\vec{x})$ having frequentist properties

Berger
Prosper (?)
Maeshima
Kruse
Linnemann
Narsky
Roe

~ coverage

mean length (median?)

$\sigma[\theta_{np}(\vec{x})]$?

- Exploits lucky data "automatically" (Berger, Roe, ?)
- Easy to combine results

$$p(\theta | \vec{x}_1, \vec{x}_2) \propto \mathcal{L}_1(\vec{x}_1 | \theta) \mathcal{L}_2(\vec{x}_2 | \theta) \pi(\theta)$$

Feldman - Cousins Issues

Feldman
Cousins
Messier
Kafka
Schnee

- Avoids unphysical & null intervals.
- Coverage (\approx independent of true θ)
- No flip-flopping
 - 2-sided interval in absence of convincing discovery
- If $n=0$, upper limit on Poisson mean decreases for increasing expected background.

G. F.
Maeshima
Conway
:

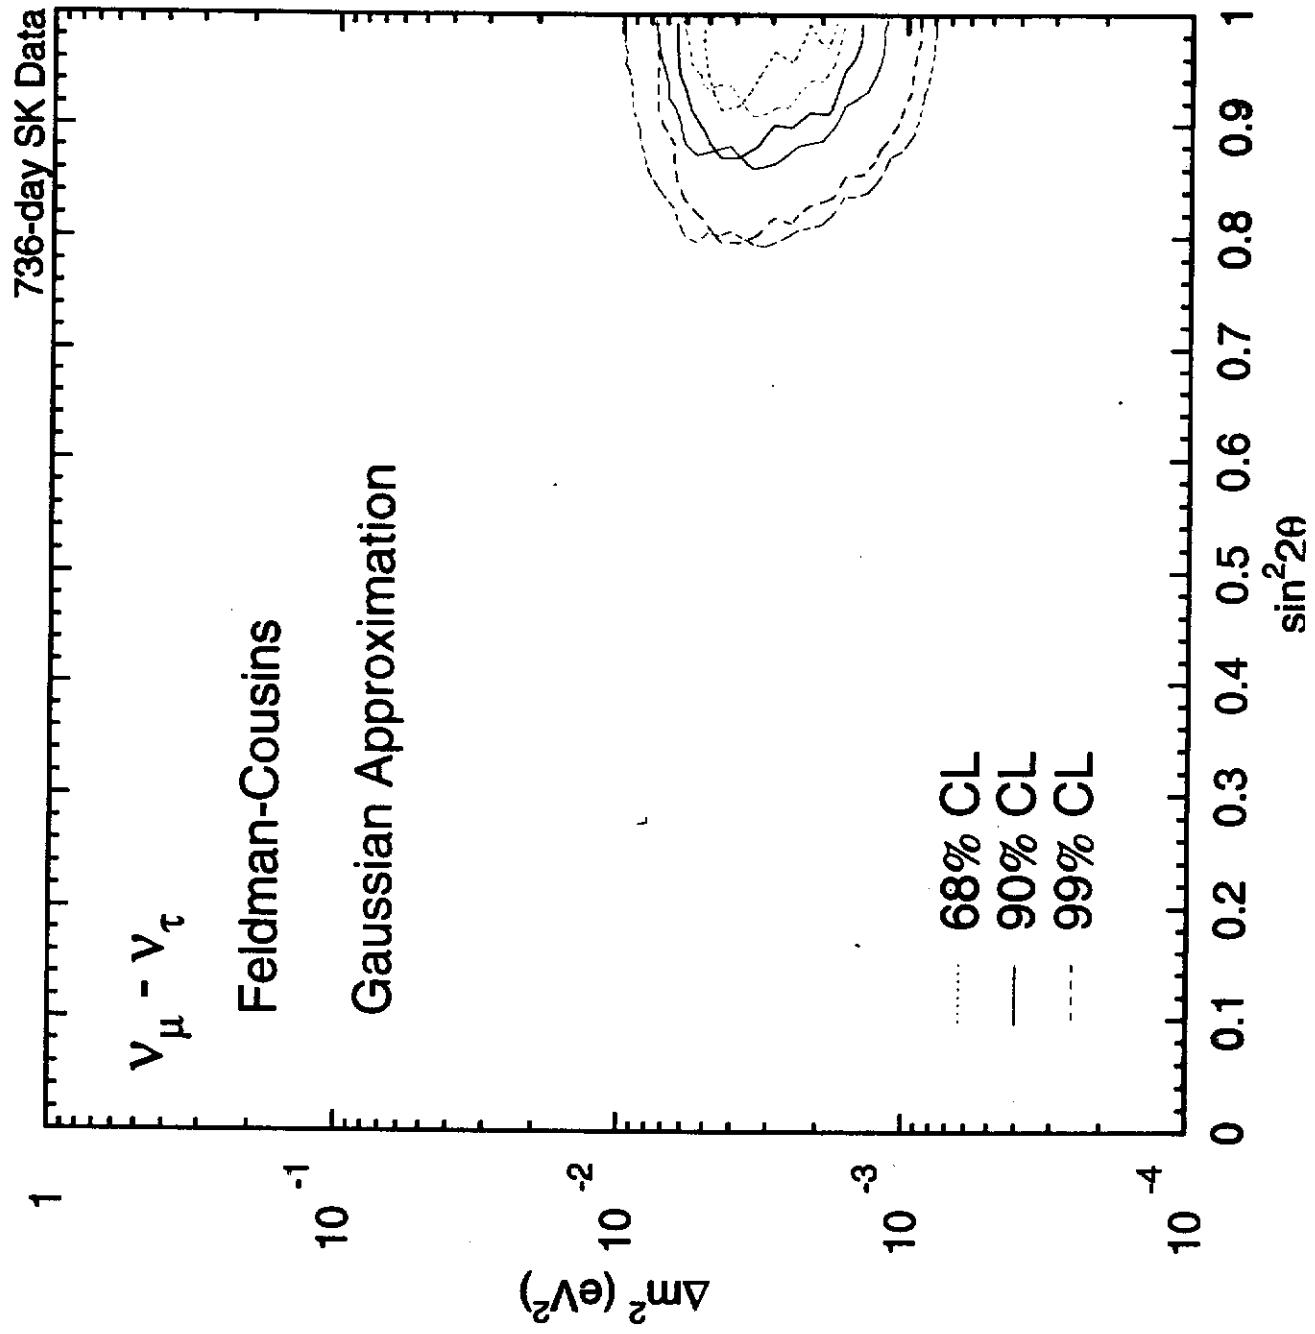
Condition on ancillary var. (Roe-Woodroffe)

Not always easy to find; (T. Berger)
doesn't always give "desired" result (Roe)

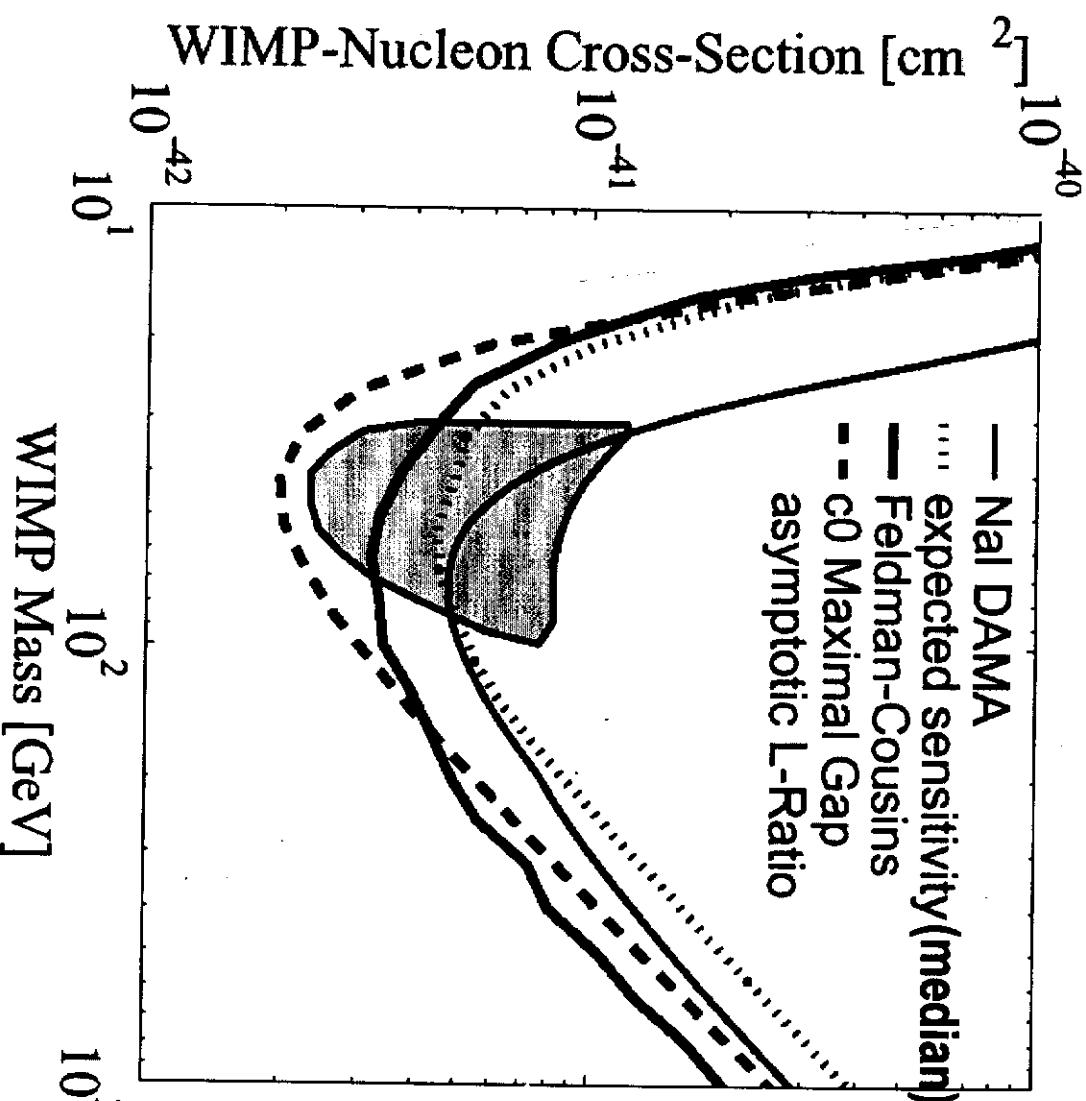
Can be computer intensive (Messier, Kafka)

What are the frequentist properties of these

Comparison of Confidence Intervals

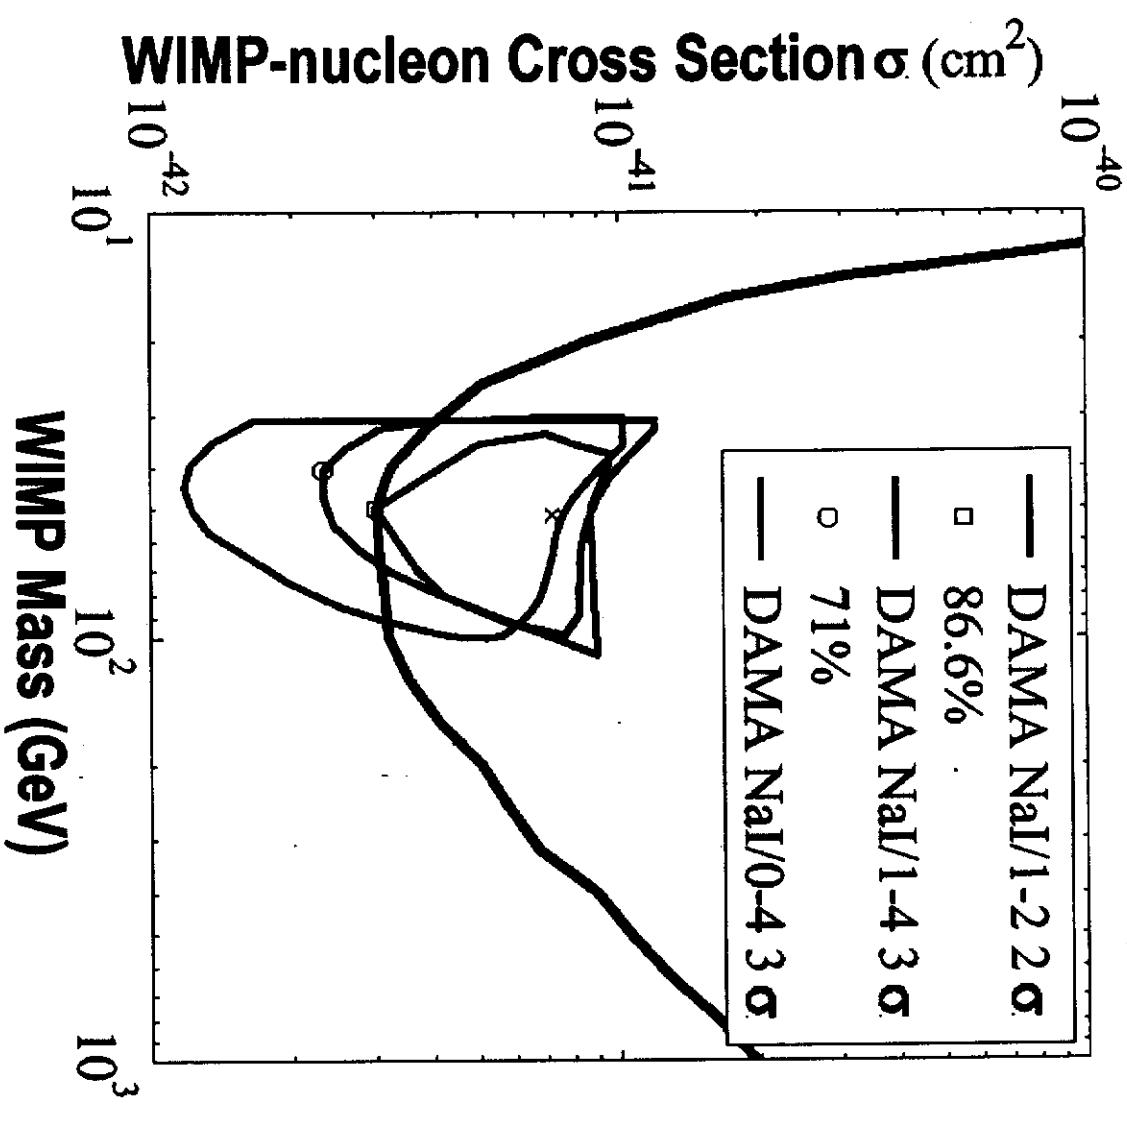


CDMS Limits ($90\% CL$)



- Because we see more multiple-scatter events than expected, limits are 50% better than expected sensitivity
- So far Bayesian method done only without energy info; results are similar to F-C.
- See <http://dmtools.berkeley.edu/limitplots/> for interactive dark matter limit plotting

Compatibility with DAMA Regions



- Bottom of DAMA NaI/1-2 2σ ($\sim 87\%$) region excluded at 86.6% CL
- Bottom of DAMA NaI/1-4 3σ ($\sim 99\%$) region excluded at 71% CL
- It does not make sense to compare to DAMA NaI/0-4 region

CL_s

Bill Murray

Prescription: $CL_s \equiv \frac{\text{P-value of } s+b}{\text{P-value of } b}$

If $CL_s < \alpha$ exclude s

\Rightarrow coverage $> 1 - \alpha$ "frequentist-safe"
(?)

Limit for $n=0$ independent of b ✓

Claim: "Outperforms objective Bayes"

(when, where etc. is this true?)

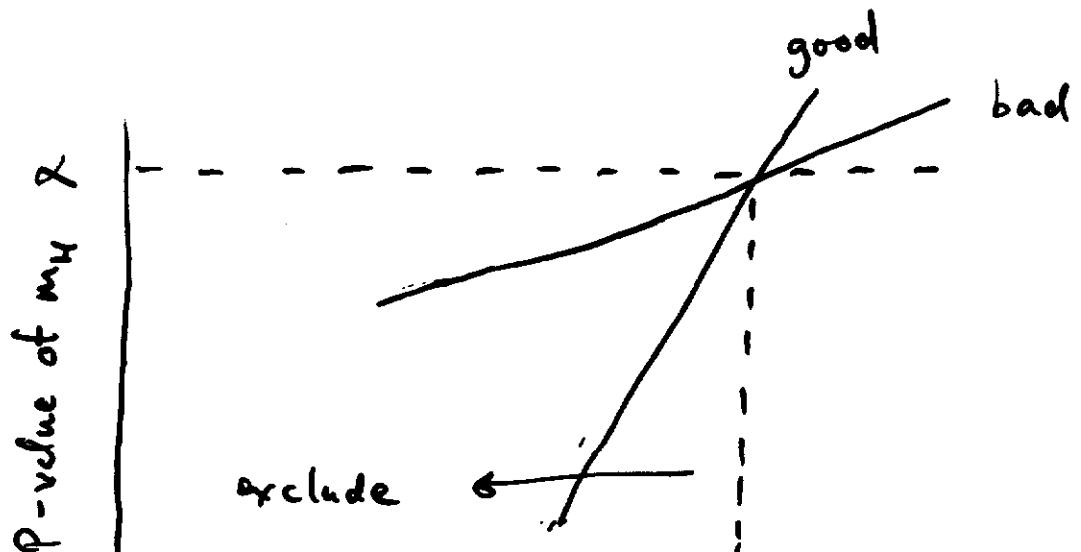
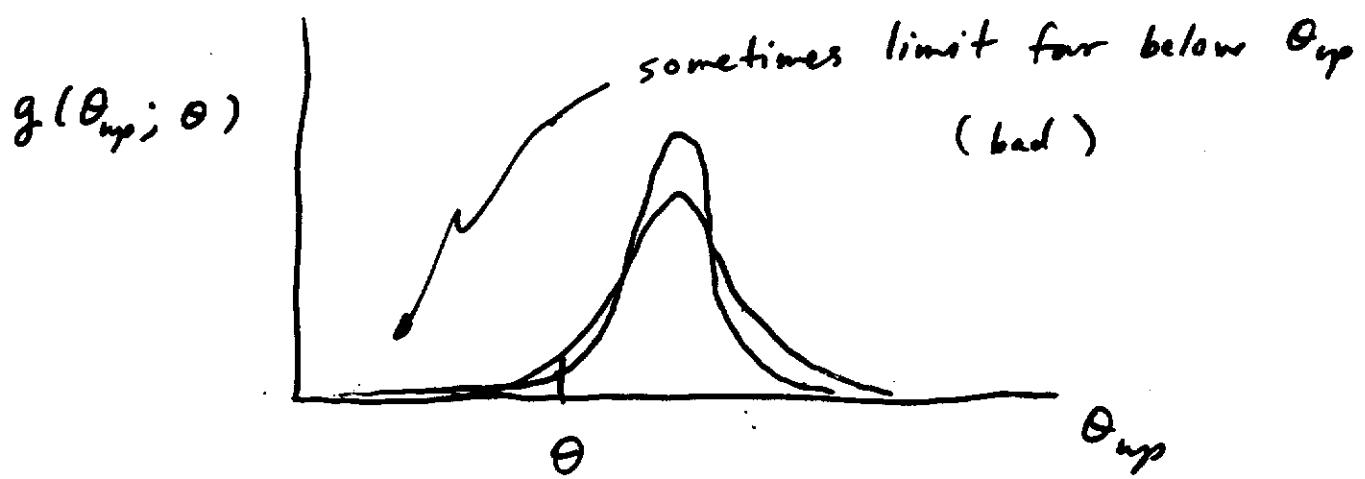
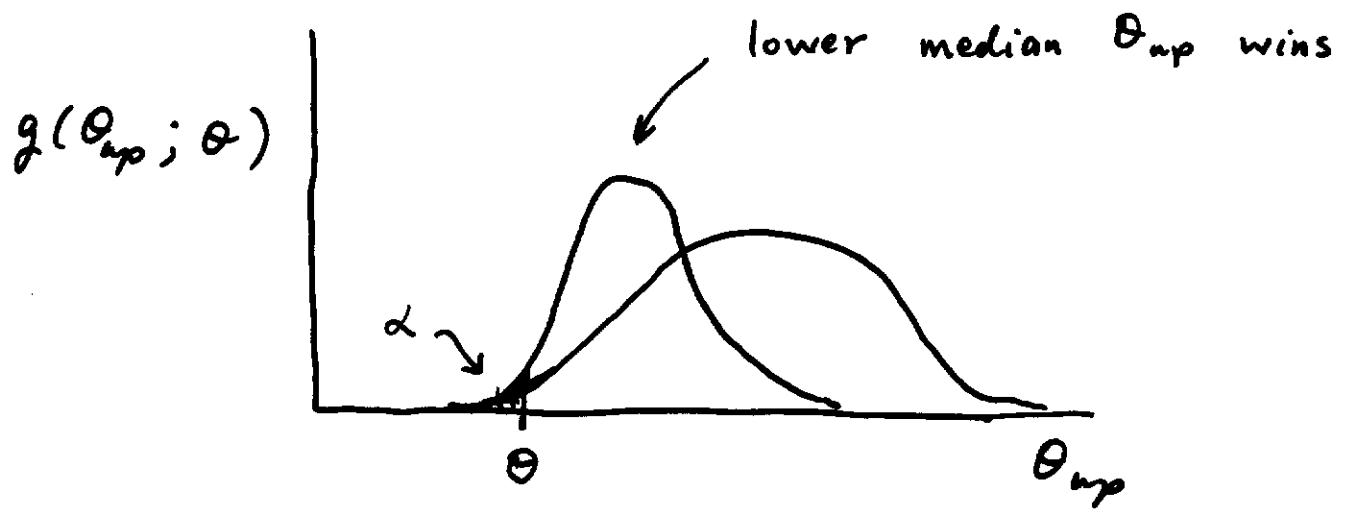
Others: "Signal Estimator method" McNamara

Hu

Nielson

:
:
:

Study distribution of limit (sampling pdf)



Systematics

Linnemann
Yellin

Bayes: nuisance parameter

$$p(\theta, b | \vec{x}) \propto \mathcal{L}(\vec{x} | \theta, b) \pi(\theta, b)$$

$$\text{e.g. } \pi(\theta, b) = \pi_b(b) \frac{1}{\sqrt{2\pi \sigma_b^2}} e^{-\frac{(b-b_0)^2}{2\sigma_b^2}}$$

Marginalize:

$$p(\theta | \vec{x}) = \int p(\theta, b | \vec{x}) db$$

F-C + Cousins - Highland \rightarrow problem with coverage?

Objective Bayes: marginalize (becomes partly subj.)

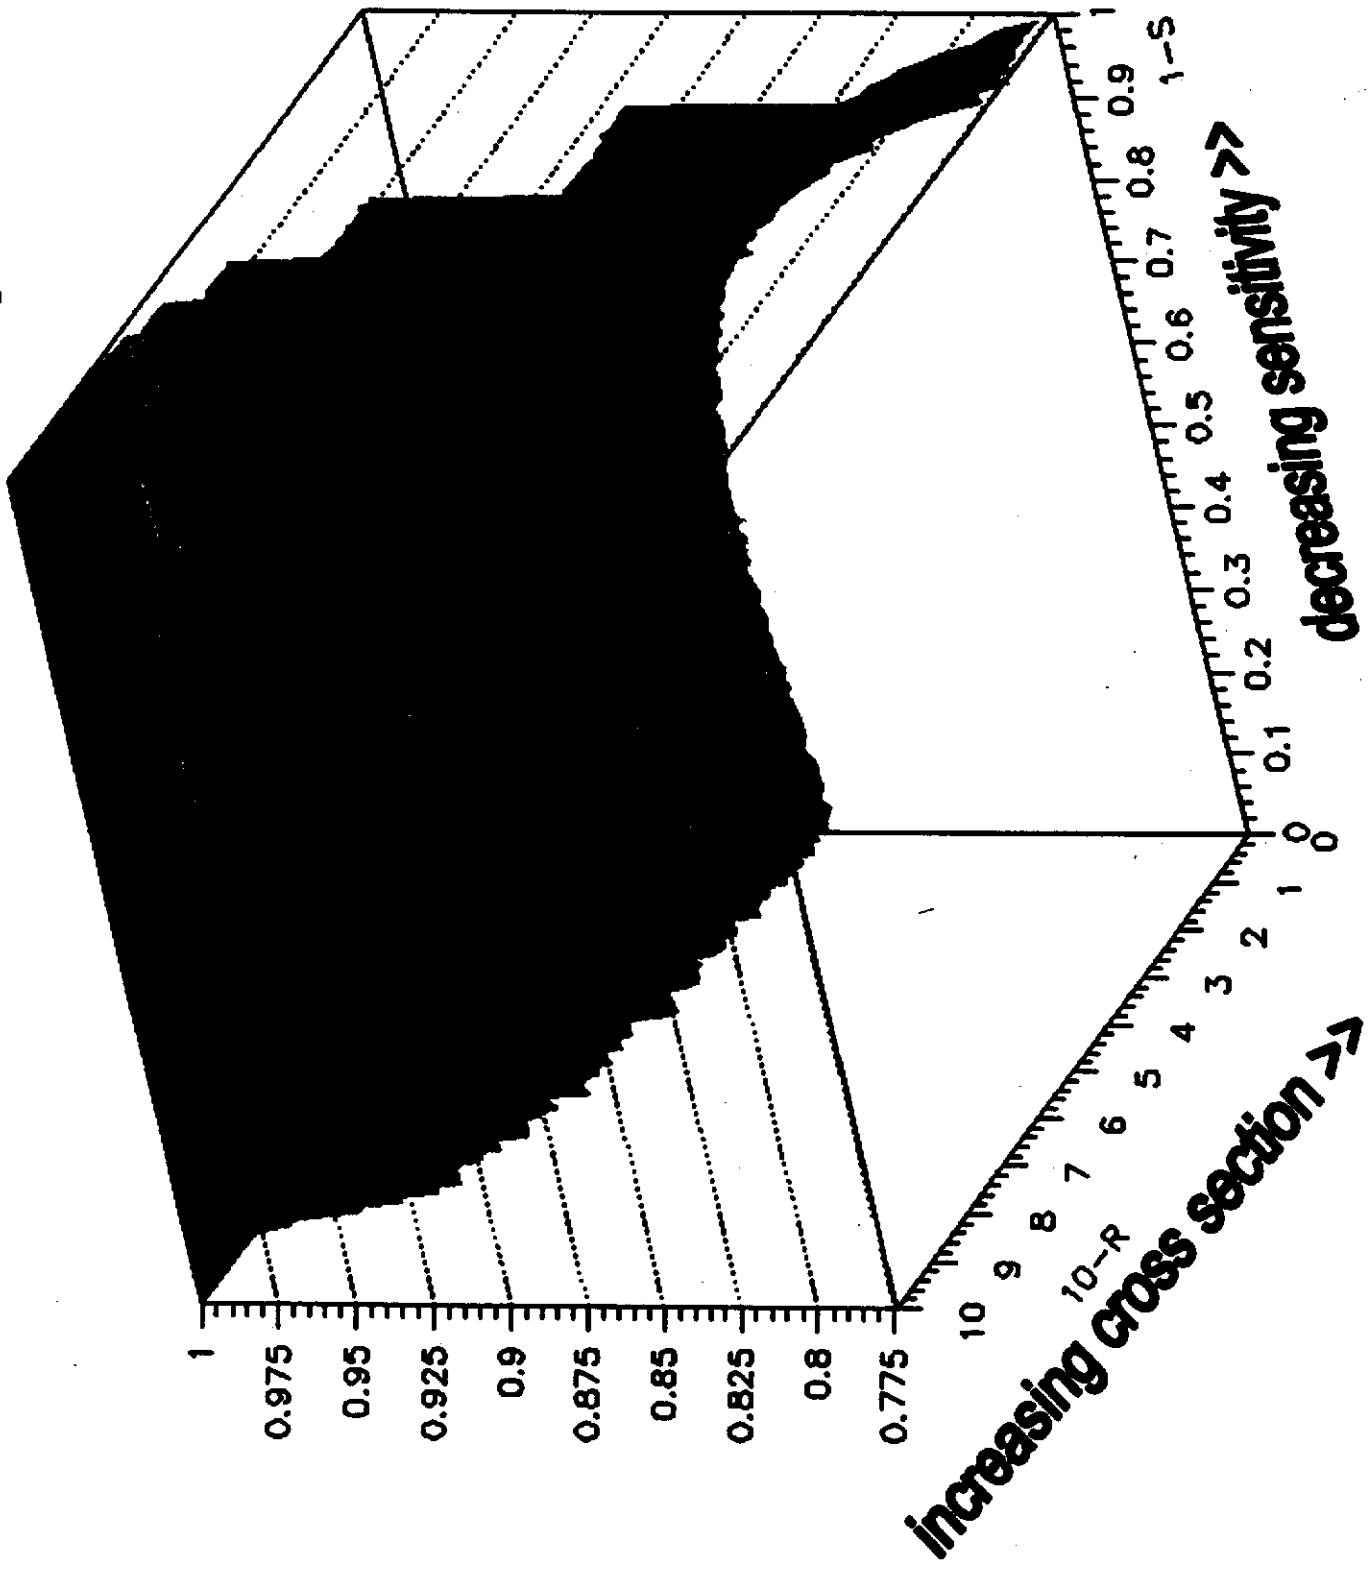
Does it make sense to report a

"marginalized \mathcal{L} " ?

$$\mathcal{L}(\theta) = \int \mathcal{L}(\theta, b) \pi_b(b) db$$

?

Coverage Probability for 90% Confidence Level for Sigma=0.25



What is the killer argument for/against

F-C + modifications

Objective Bayes

CL,

?

Which limit best allows the consumer
to calculate

$$p(\theta | \text{limit}) \propto p(\text{limit} | \theta) \pi(\theta)$$

consumer's prior

so that it gives the most (clearest?)
information about θ where it counts?

N.B. limit alone not enough; need

$$p(\text{limit} | \theta) \rightarrow E[\text{limit}]$$

$$\sigma[\text{limit}] \quad (\text{Raja.})$$

(Recall serious Bayesian consumer will want

$$p(\theta | \tilde{x}) \propto \lambda(\tilde{x} | \theta) \pi(\theta)$$

\Rightarrow limit for "casual Bayesians".)

Expected Confidence Levels

What does expected mean?

- Mean

- Has normally been used by us.

- Median

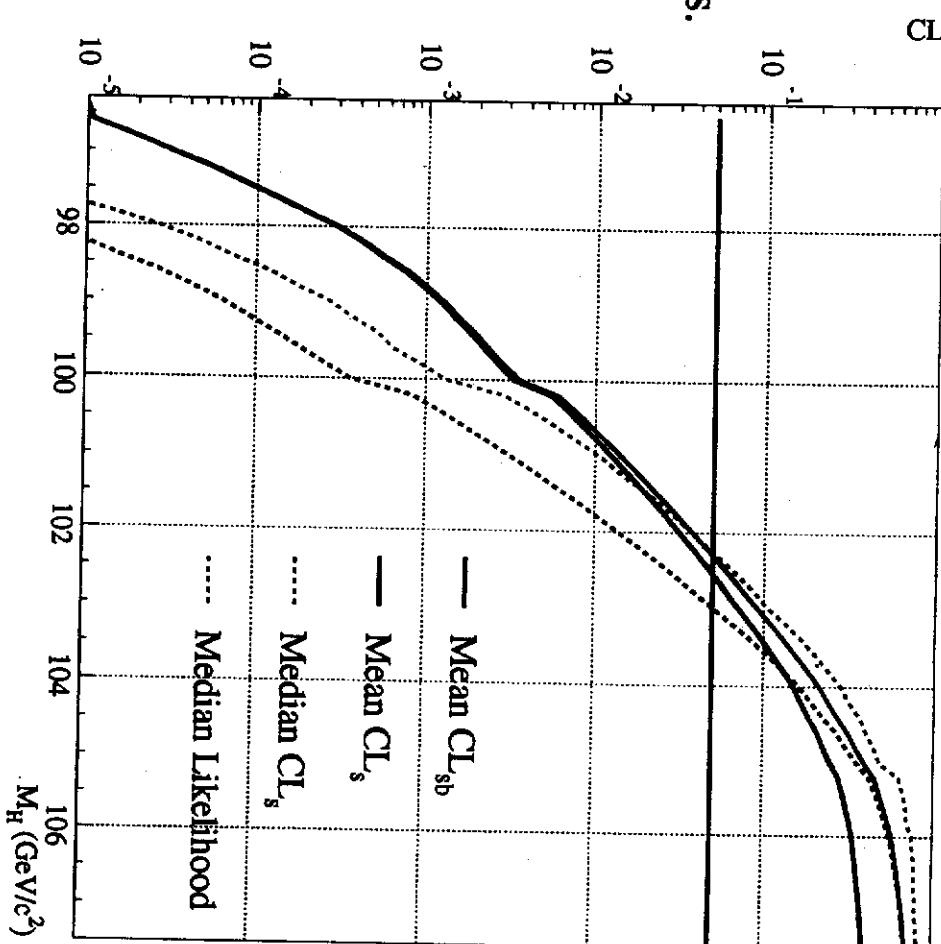
- No dependence on metric

Careful: Both are used here!

Expected limits:

$CL_s \cdot 3\text{GeV}$ below CL_{sb}

LR: 1GeV below CL_{sb}

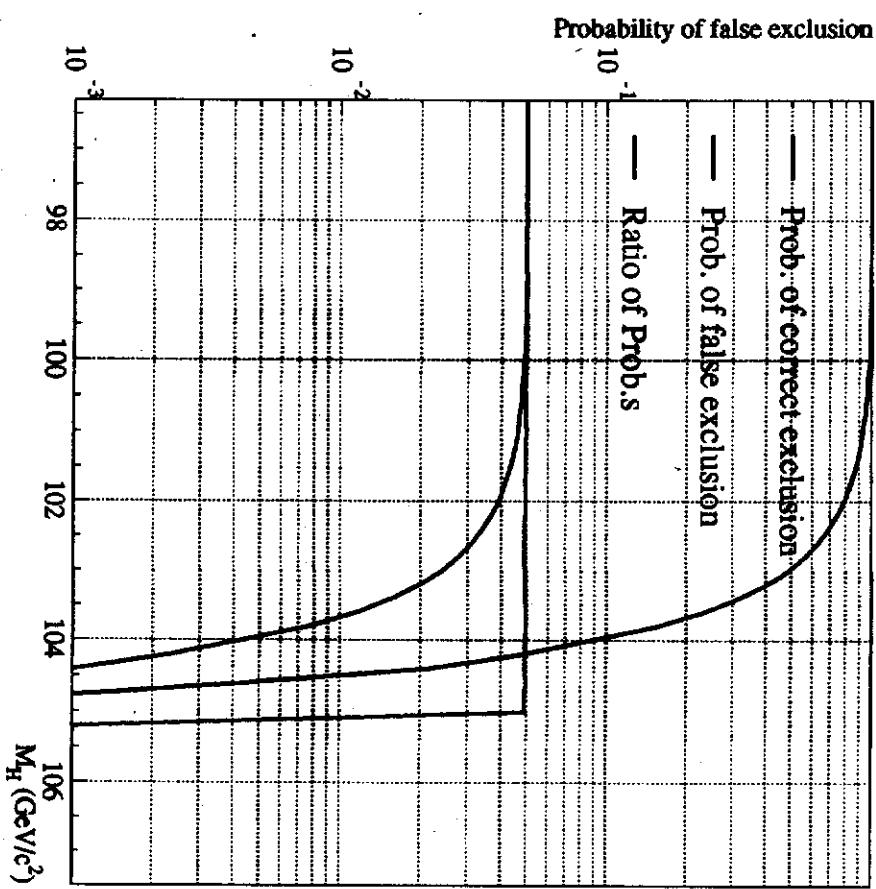


Probability of Exclusion

Define exclusion as CLs<0.05

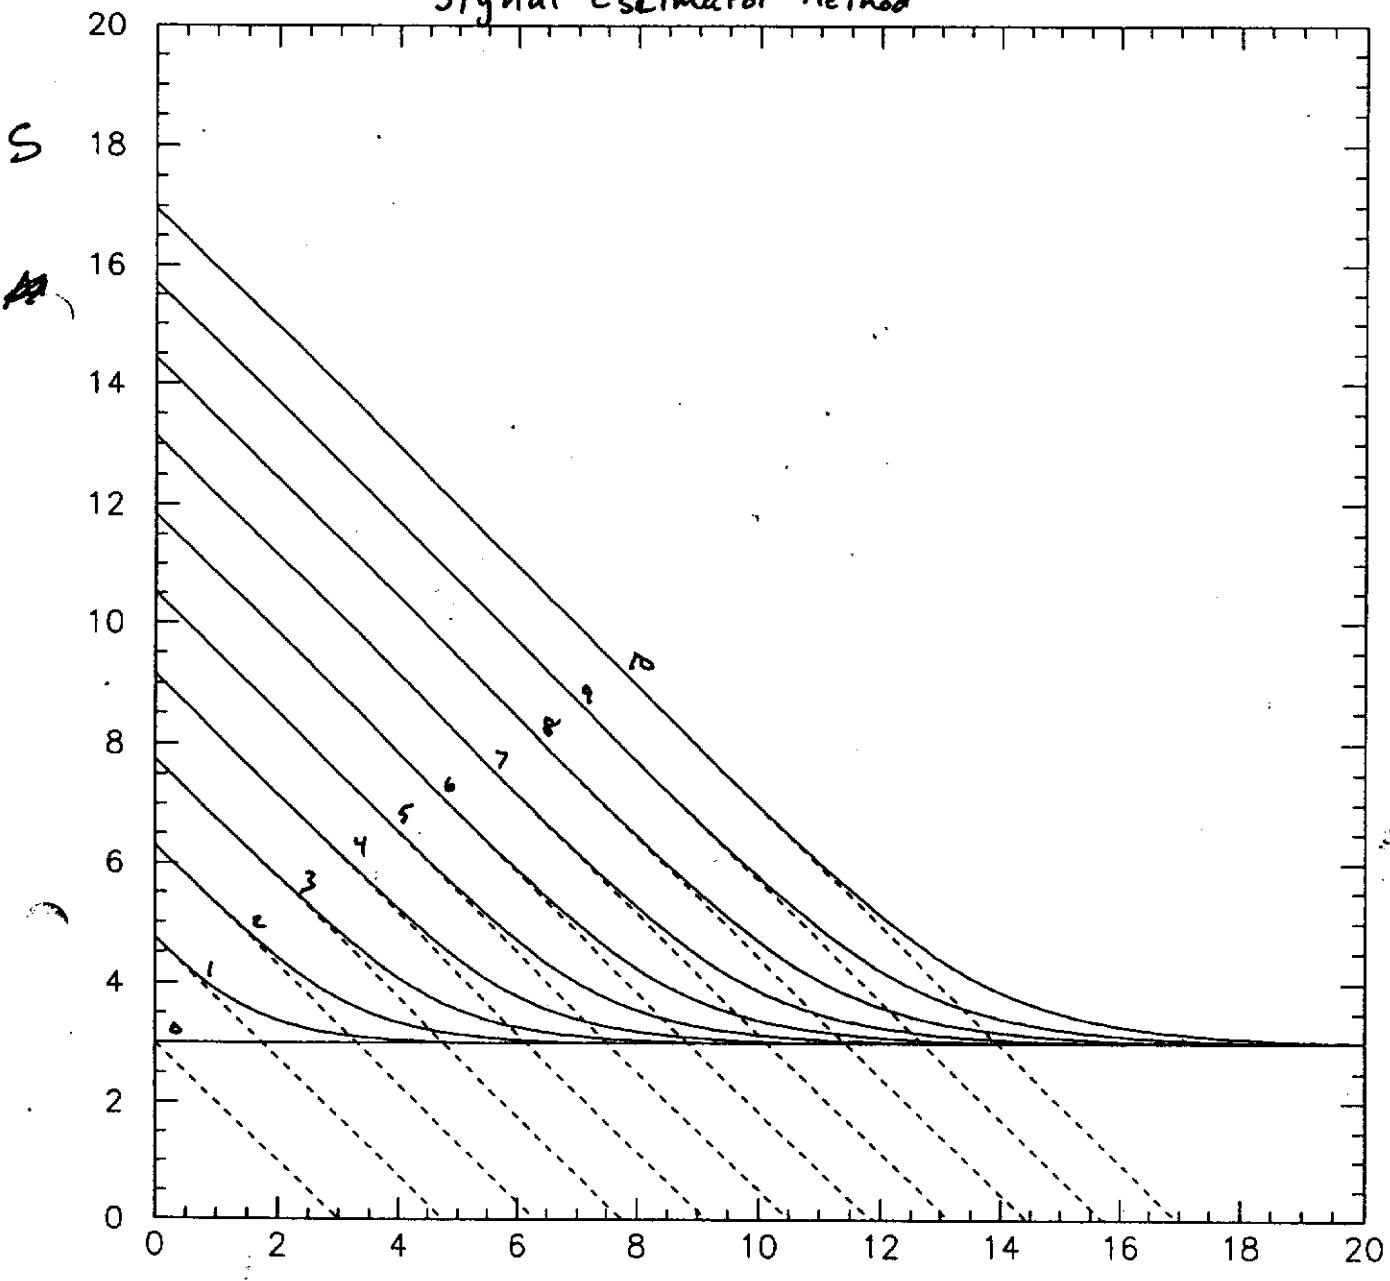
Probability of false exclusion
should be 5% - but is less

Significant overcoverage

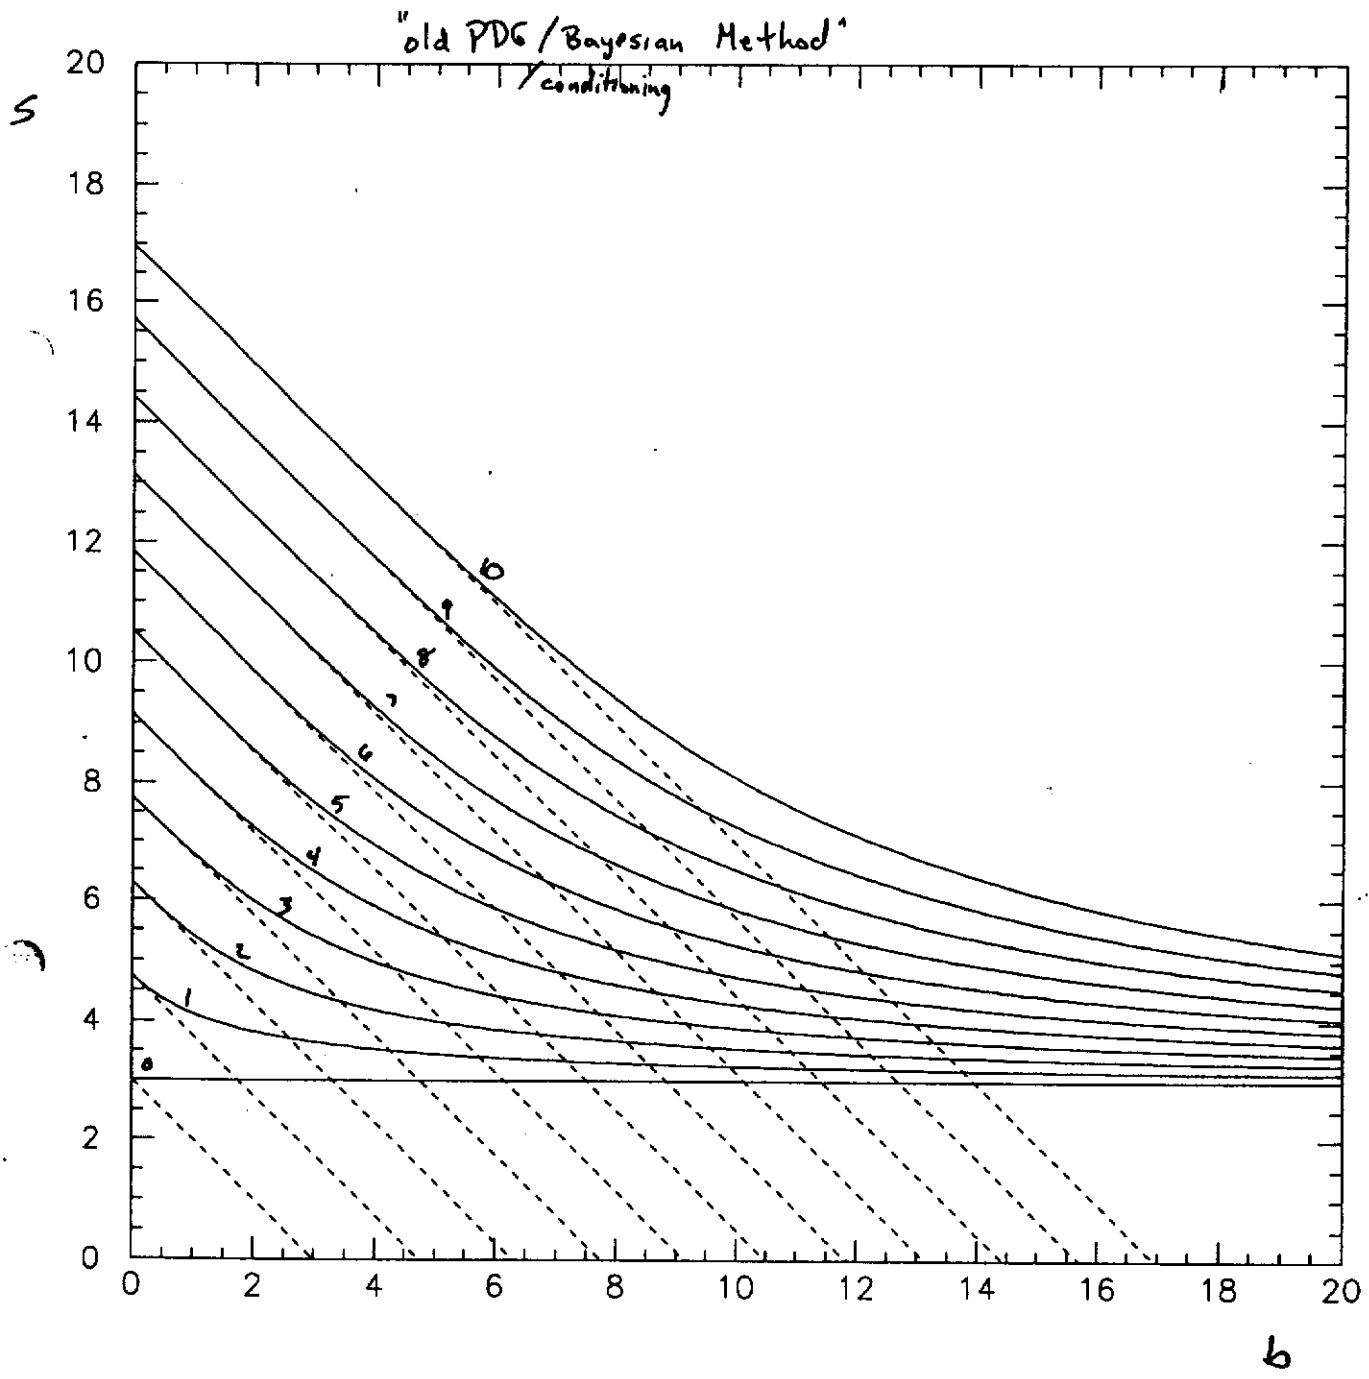


False exclusion
rate is always 5%
*of the true
exclusion rate*

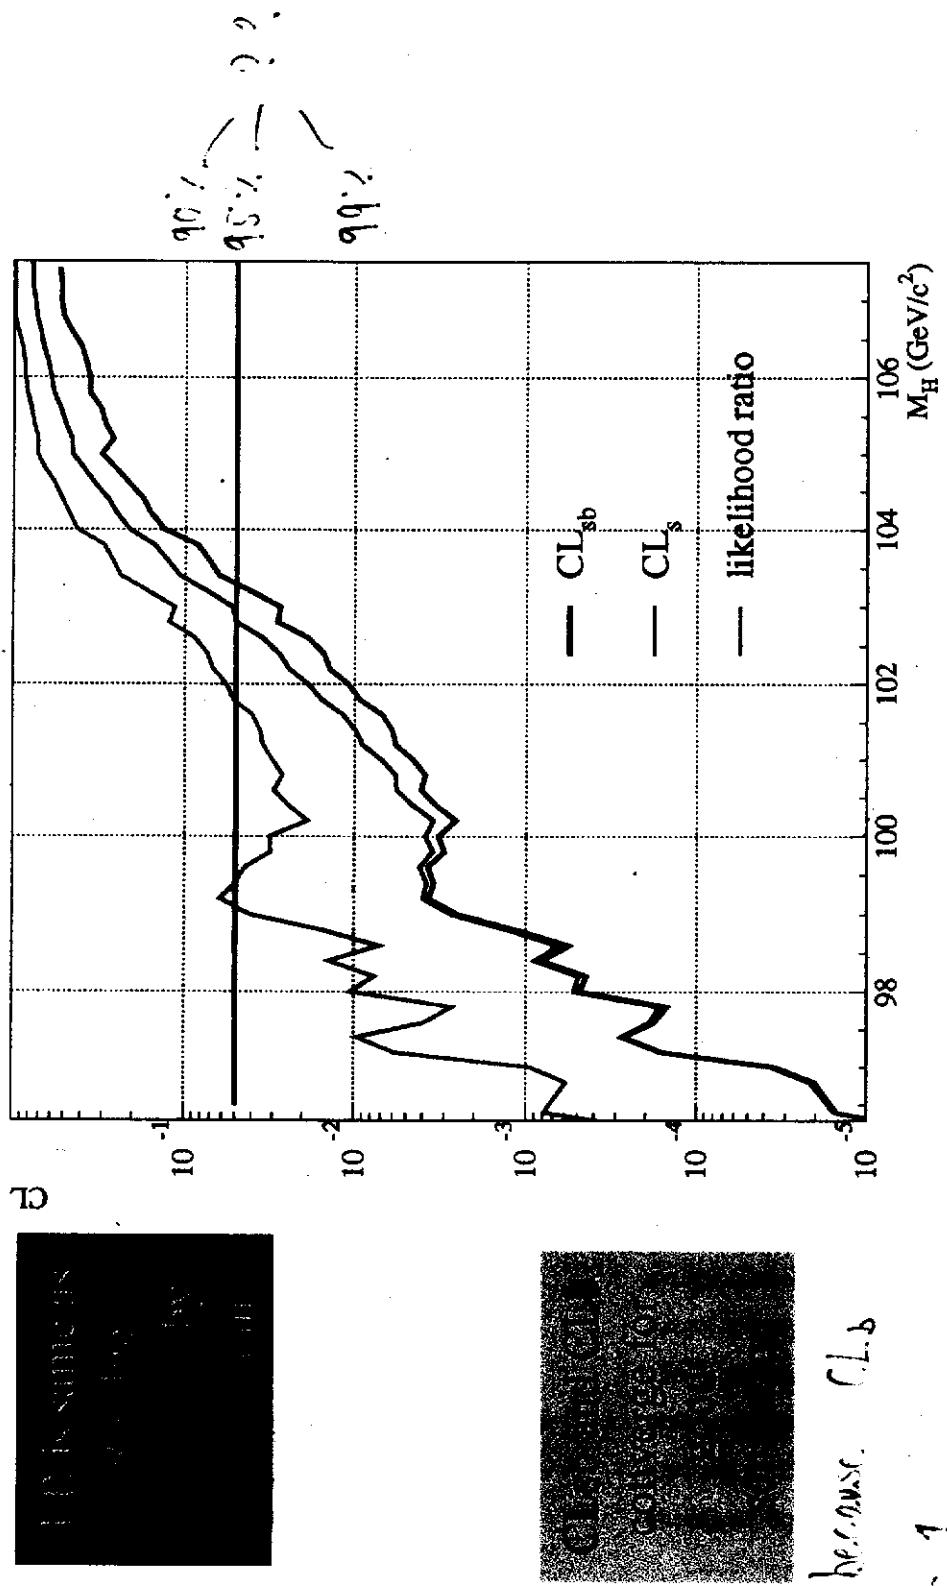
"Signal Estimator Method"



b



Observed Confidence Levels



Open Questions, etc.

- Is the stopping rule problem a non-problem?
- Does a Bayesian analysis prevent one from addressing goodness-of-fit?
- Should 95% coverage mean "minimum" or "typical"? (Why 95%?)
- Subjective \neq arbitrary

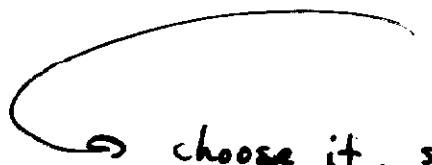
Elements of arbitrariness:

Bayes - choice of prior;

"metric" ($\theta \rightarrow h(\theta)$),

Freq. - choice of method used to compactify result (estimator);

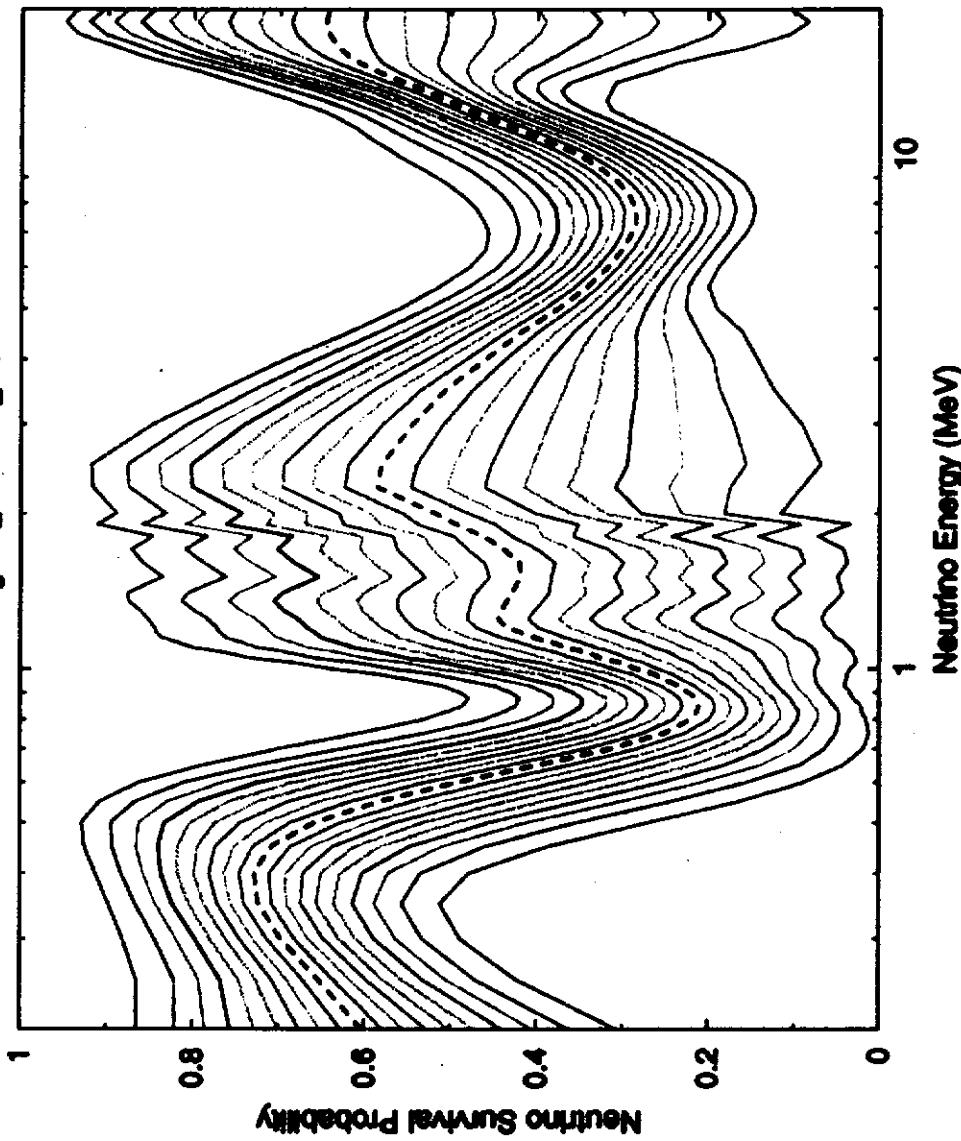
choice of ensemble.



choose it, state it, use it, move on

$\Pr(p|D)$: Active Neutrinos

Survival Probability vs Neutrino Energy for Parametric Method
Assuming $\nu_{\text{e}} \rightarrow \nu_{\mu}$



For the "Results" section:

- Avoid subjectivity with respect to preference of theory (parameter)
- Summarize results:
 - compact
 - maximize useful information

→ functions of the data with well defined properties given a theory.

likelihood function $\mathcal{L} = P(\text{data} | \text{theory})$

limits

P-values (CL) vs. parameter (goodness-of-fit statistics)
- Provide input for (Bayesian) consumer who will want consumer's prior $p(\text{theory} | \text{result}) \propto p(\text{result} | \text{theory}) \pi(\text{theory})$

e.g. $p(m_H | \text{your limit}) \propto p(\text{your limit} | m_H) \pi(m_H)$

- coverage : like it or hate it ?
- ability to combine results
(→ x ?)
- digestable by non-specialists,
e.g. your ed. board
- do we need a consensus ?
- Future workshop: systematics